
Chapter1 Linear Inverse Problems with Norm and Sparsity Constraints

Volkan Cevher

volkan.cevher@epfl.ch

Laboratory for Information and Inference Systems

Ecole Polytechnique Federale de Lausanne

Sina Jafarpour

sina2jp@yahoo-inc.com

Multimedia Research Group

Yahoo! Research

Anastasios Kyrillidis

anastasios.kyrillidis@epfl.ch

Laboratory for Information and Inference Systems

Ecole Polytechnique Federale de Lausanne

*

We describe two nonconventional algorithms for linear regression, called GAME and CLASH. The salient characteristics of these approaches is that they exploit the convex ℓ_1 -ball and non-convex ℓ_0 -sparsity constraints jointly in sparse recovery. To establish the theoretical approximation guarantees of GAME and CLASH, we cover an interesting range of topics from game theory, convex and combinatorial optimization. We illustrate that these approaches lead to improved theoretical guarantees and empirical performance beyond convex and non-convex solvers alone.

1.1 Introduction

Sparse approximation is a fundamental problem in compressed sensing [1, 2], as well as in many other signal processing and machine learning applications including variable selection in regression [3, 4, 5], graphical model selection [6, 7], and sparse principal component analysis [8, 9]. In sparse approximation, one is provided with

*, Authors are in alphabetical order.

a dimensionality reducing measurement matrix $\Phi \in \mathbb{R}^{M \times N}$ ($M < N$), and a low dimensional vector $\mathbf{f} \in \mathbb{R}^M$ such that:

$$\mathbf{f} = \Phi \boldsymbol{\alpha}^* + \mathbf{n}, \quad (1.1)$$

where $\boldsymbol{\alpha}^* \in \mathbb{R}^N$ is the high-dimensional signal of interest and $\mathbf{n} \in \mathbb{R}^M$ is a potential additive noise term with $\|\mathbf{n}\|_2 \leq \sigma$.

In this work, we assume $\boldsymbol{\alpha}^*$ is a k -sparse signal or is sufficiently approximated by a k -sparse vector. The goal of sparse approximation algorithms is then to find a sparse vector $\hat{\boldsymbol{\alpha}} \in \mathbb{R}^N$ such that $\Phi \hat{\boldsymbol{\alpha}} - \mathbf{f}$ is small in an appropriate norm. In this setting, the ℓ_0 -minimization problem emerges naturally as a suitable solver to recover $\boldsymbol{\alpha}^*$ in (1.1):

$$\underset{\boldsymbol{\alpha} \in \mathbb{R}^N}{\text{minimize}} \quad \|\boldsymbol{\alpha}\|_0 \quad \text{subject to} \quad \|\mathbf{f} - \Phi \boldsymbol{\alpha}\|_2 \leq \sigma, \quad (1.2)$$

where $\|\boldsymbol{\alpha}\|_0$ counts the nonzero elements (the sparsity) of $\boldsymbol{\alpha}$.

Unfortunately, solving (1.2) is a challenging task with exponential time complexity. Representing the set of all k -sparse vectors as:

$$\Delta_{\ell_0}(k) \doteq \{\boldsymbol{\alpha} \in \mathbb{R}^N : \|\boldsymbol{\alpha}\|_0 \leq k\}, \quad (1.3)$$

hard thresholding algorithms [10, 11, 12, 13, 14] abandon this approach in favor of greedy selection where a putative k -sparse solution is iteratively refined using local decision rules. To this end, hard thresholding methods consider the following ℓ_0 -constrained least squares problem formulation as an alternative to (1.2):

$$\underset{\boldsymbol{\alpha} \in \mathbb{R}^N}{\text{minimize}} \quad \|\mathbf{f} - \Phi \boldsymbol{\alpha}\|_2^2 \quad \text{subject to} \quad \boldsymbol{\alpha} \in \Delta_{\ell_0}(k). \quad (1.4)$$

These methods feature computational advantages and also are backed up with a great deal of theory for estimation guarantees.

In contrast, convex optimization approaches change the problem formulations above by “convexifying” the combinatorial ℓ_0 -constraint with the sparsity inducing convex ℓ_1 -norm.¹ As a result, (1.2) is transformed into the ℓ_1 -minimization, also known as the Basis Pursuit (BP) problem [15]:

$$\underset{\boldsymbol{\alpha} \in \mathbb{R}^N}{\text{minimize}} \quad \|\boldsymbol{\alpha}\|_1 \quad \text{subject to} \quad \|\mathbf{f} - \Phi \boldsymbol{\alpha}\|_2 \leq \sigma. \quad (1.5)$$

Similarly, the famous Lasso algorithm [16] can be considered as a relaxation of (1.4):

$$\underset{\boldsymbol{\alpha} \in \mathbb{R}^N}{\text{minimize}} \quad \|\mathbf{f} - \Phi \boldsymbol{\alpha}\|_2^2 \quad \text{subject to} \quad \boldsymbol{\alpha} \in \Delta_{\ell_1}(\tau), \quad (1.6)$$

where $\Delta_{\ell_1}(\tau)$ is the set of all vectors inside the hyper-diamond of radius τ :

$$\Delta_{\ell_1}(\tau) \doteq \{\boldsymbol{\alpha} \in \mathbb{R}^N : \|\boldsymbol{\alpha}\|_1 \leq \tau\}. \quad (1.7)$$

1. Note that this is not a true convexification, since the ℓ_0 -ball does not have a scale.

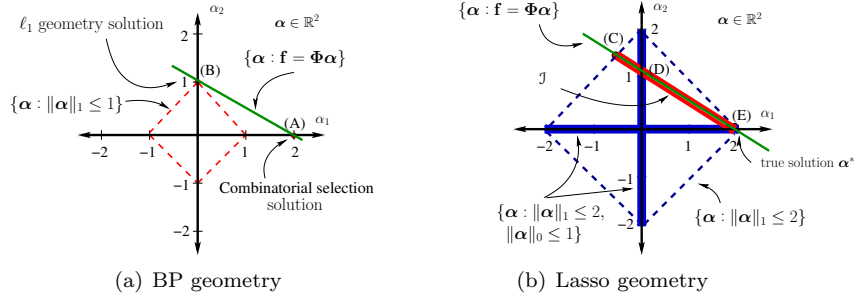


Figure 1.1: Geometric interpretation of the selection process for a simple test case $\mathbf{f} = \Phi\boldsymbol{\alpha}^*$ where $\|\boldsymbol{\alpha}^*\|_0 = 1$.

While both convex and non-convex problem formulations can find the true problem solution under various theoretical assumptions, one can easily find examples in practice where either one can fail. Borrowing from [17], we provide an illustrative example in \mathbb{R}^2 for the noiseless case in Fig. 1.1. In (1.2), combinatorial-based approaches can identify the admissible set of 1-sparse solutions. If a greedy selection rule is used to arbitrate these solutions, then such an approach could pick (A). In contrast, the BP algorithm selects a solution (B), and misses the candidate solution (A) as it cannot exploit prior knowledge concerning the discrete structure of $\boldsymbol{\alpha}^*$.

To motivate our discussion in this book chapter, let us assume that we have the *true model parameters* $\|\boldsymbol{\alpha}^*\|_0 = k$ and $\|\boldsymbol{\alpha}^*\|_1 = \tau$. Let us then consider geometrically the—unfortunate but common—case where the kernel of Φ , $\ker(\Phi)$, intersects with the tangent cone $T_{\|\boldsymbol{\alpha}\|_1 \leq \tau}(\boldsymbol{\alpha}^*) = \{s(\mathbf{y} - \boldsymbol{\alpha}^*) : \|\mathbf{y}\|_1 \leq \tau \text{ and } s \geq 0\}$ at the true vector $\boldsymbol{\alpha}^*$ (cf., (E) in Fig. 1.1(b)). From the Lasso perspective, we are stuck with the large continuum of solutions based on the geometry, as described by the set $\mathcal{J} = \ker(\Phi) \cap T_{\|\boldsymbol{\alpha}\|_1 \leq \tau}(\boldsymbol{\alpha}^*)$, as illustrated in Figure 1.1(b) within the box.

Without further information about the discrete nature of $\boldsymbol{\alpha}^*$, a convex optimization algorithm solving the Lasso problem can arbitrarily select a vector from \mathcal{J} . By forcing basic solutions in optimization, we can reduce the size of the solution space to $\mathcal{L} = \mathcal{J} \cap \{\|\boldsymbol{\alpha}\|_1 = 1\}$, which is constituted by the sparse vectors (C) and (E). Note that \mathcal{L} might be still large in high dimensions. However, in this scenario, adding the $\Delta_{\ell_0}(k)$ constraints, we can make precise selections (e.g., exactly 1-sparse), significantly reduce the candidate solution set, and, in many cases, can obtain the correct solution (E) if we leverage the norm constraint.

Contents of this book chapter: Within this context, we describe two efficient, sparse approximation algorithms, called GAME and CLASH, that operate over sparsity and ℓ_1 -norm constraints. They address the following nonconvex problem:

$$\underset{\boldsymbol{\alpha} \in \Delta_{\ell_0, \ell_1}(k, \tau)}{\text{minimize}} \quad \|\Phi\boldsymbol{\alpha} - \mathbf{f}\|_q, \quad (1.8)$$

where $\Delta_{\ell_1}(\tau)$ is the set of all k -sparse vectors in $\Delta_{\ell_1}(\tau)$:

$$\Delta_{\ell_0, \ell_1}(k, \tau) \doteq \{\boldsymbol{\alpha} \in \mathbb{R}^N : \|\boldsymbol{\alpha}\|_0 \leq k \text{ and } \|\boldsymbol{\alpha}\|_1 \leq \tau\}. \quad (1.9)$$

To introduce the *Game-theoretic Approximate Matching Estimator* (GAME) method, we reformulate (1.8) as a zero-sum game. GAME then efficiently obtains a sparse approximation for the optimal game solution. GAME employs a primal-dual scheme, and require $\tilde{O}(k)$ iterations in order to find a k -sparse vector with $O(k^{-0.5})$ additive approximation error.

To introduce the *Combinatorial selection and Least Absolute SHrinkage* operator CLASH, we recall hard thresholding methods and explain how to incorporate the ℓ_1 norm constraint. A key feature of the CLASH approach is that it allows us to exploit ideas from the model-based compressive sensing (model-CS) approach, where selections can be driven by a structured sparsity model [18, 19].

We emphasize again that since $\Delta_{\ell_0, \ell_1}(k, \tau)$ is not convex, the optimization problem (1.8) is not a convex optimization problem. However, we can still derive theoretical approximation guarantees of both algorithms. For instance, we can prove that for every dimension reducing matrix Φ , and every measurement vector \mathbf{f} , GAME can find a vector $\hat{\alpha} \in \Delta_{\ell_0, \ell_1}(k, \tau)$ with

$$\|\Phi \hat{\alpha} - \mathbf{f}\|_q \leq \min_{\alpha \in \Delta_{\ell_0, \ell_1}(k, \tau)} \|\Phi \alpha - \mathbf{f}\|_q + \tilde{O}\left(\frac{1}{\sqrt{k}}\right), \quad (1.10)$$

where q is a positive integer. This sparse approximation framework surprisingly works for any matrix Φ . Compared to the GAME algorithm, CLASH requires stronger assumptions on the measurement matrix for estimation guarantees. However, these assumptions, in the end, lead to improved empirical performance.

1.2 Preliminaries

Here, we cover basic mathematical background that is used in establishing algorithmic guarantees in the sequel.

1.2.1 Bregman Projections

Bregman divergences or Bregman distances are an important family of distances that all share similar properties [20, 21].

Definition 1.1 (Bregman Distance). *Let $\mathcal{R} : \mathcal{S} \rightarrow \mathbb{R}$ be a continuously-differentiable real-valued and strictly convex function defined on a closed convex set \mathcal{S} . The Bregman distance associated with \mathcal{R} for points \mathbf{P} and \mathbf{Q} is:*

$$\mathcal{B}_{\mathcal{R}}(\mathbf{P}, \mathbf{Q}) = \mathcal{R}(\mathbf{P}) - \mathcal{R}(\mathbf{Q}) - \langle (\mathbf{P} - \mathbf{Q}), \nabla \mathcal{R}(\mathbf{Q}) \rangle.$$

Table 1.1 summarizes examples of the most widely used Bregman functions and the corresponding Bregman distances.

The Bregman distance has several important properties that we will use later in analyzing our sparse approximation algorithm.




Figure 1.2: The Bregman divergence associated with a continuously-differentiable real-valued and strictly convex function \mathcal{R} is the vertical distance at \mathbf{P} between the graph of \mathcal{R} and the line tangent to the graph of \mathcal{R} in \mathbf{Q} .

Table 1.1: Summary of the most popular Bregman functions and their corresponding Bregman distances. Here Φ is a positive semidefinite matrix.

Name	Bregman Function ($\mathcal{R}(\mathbf{P})$)	Bregman Distance ($\mathcal{B}_{\mathcal{R}}(\mathbf{P}, \mathbf{Q})$)
Squared Euclidean	$\ \mathbf{P}\ _2^2$	$\ \mathbf{P} - \mathbf{Q}\ _2^2$
Squared Mahalanobis	$\langle \mathbf{P}, \Phi \mathbf{P} \rangle$	$\langle (\mathbf{P} - \mathbf{Q}), \Phi (\mathbf{P} - \mathbf{Q}) \rangle$
Entropy	$\sum_i P_i \log P_i - P_i$	$\sum_i P_i \log \frac{P_i}{Q_i} - \sum_i (P_i - Q_i)$
Itakura-Saito	$\sum_i -\log P_i$	$\sum_i \left(\frac{P_i}{Q_i} - \log \frac{P_i}{Q_i} + 1 \right)$

Theorem 1.2. *Bregman distance satisfies the following properties:*

- (P1). $\mathcal{B}_{\mathcal{R}}(\mathbf{P}, \mathbf{Q}) \geq 0$, and the equality holds if and only if $\mathbf{P} = \mathbf{Q}$.
- (P2). For every fixed \mathbf{Q} if we define $\mathcal{G}(\mathbf{P}) = \mathcal{B}_{\mathcal{R}}(\mathbf{P}, \mathbf{Q})$, then

$$\nabla \mathcal{G}(\mathbf{P}) = \nabla \mathcal{R}(\mathbf{P}) - \nabla \mathcal{R}(\mathbf{Q}).$$

- (P3). Three point property: For every \mathbf{P}, \mathbf{Q} and \mathbf{T} in \mathcal{S}

$$\mathcal{B}_{\mathcal{R}}(\mathbf{P}, \mathbf{Q}) = \mathcal{B}_{\mathcal{R}}(\mathbf{P}, \mathbf{T}) + \mathcal{B}_{\mathcal{R}}(\mathbf{T}, \mathbf{Q}) + \langle (\mathbf{P} - \mathbf{T}), \nabla \mathcal{R}(\mathbf{Q}) - \nabla \mathcal{R}(\mathbf{T}) \rangle.$$

- (P4). For every $\mathbf{P}, \mathbf{Q} \in \mathcal{S}$,

$$\mathcal{B}_{\mathcal{R}}(\mathbf{P}, \mathbf{Q}) + \mathcal{B}_{\mathcal{R}}(\mathbf{Q}, \mathbf{P}) = \langle (\mathbf{P} - \mathbf{Q}), (\nabla \mathcal{R}(\mathbf{P}) - \nabla \mathcal{R}(\mathbf{Q})) \rangle.$$

Proof. All four properties follow directly from Definition 1.1. \square

Now that we are equipped with the properties of Bregman distances, we are ready to define Bregman projections of points into convex sets.

Definition 1.3 (Bregman Projection). *Let $\mathcal{R} : \mathcal{S} \rightarrow \mathbb{R}$ be a continuously-differentiable real-valued and strictly convex function defined on a closed convex set \mathcal{S} . Let Ω be a closed subset of \mathcal{S} . Then, for every point \mathbf{Q} in \mathcal{S} , the Bregman projection of \mathbf{Q} into Ω , denoted as $\mathcal{P}_{\Omega}(\mathbf{Q})$ is*

$$\mathcal{P}_{\Omega}(\mathbf{Q}) \doteq \arg \min_{\mathbf{P} \in \Omega} \mathcal{B}_{\mathcal{R}}(\mathbf{P}, \mathbf{Q}).$$

Bregman projections satisfy a generalized Pythagorean Theorem.

Theorem 1.4 (Generalized Pythagorean Theorem [20]). *Let $\mathcal{R} : \mathcal{S} \rightarrow \mathbb{R}$ be a continuously-differentiable real-valued and strictly convex function defined on a closed convex set \mathcal{S} . Let Ω be a closed subset of \mathcal{S} . Then for every $\mathbf{P} \in \Omega$ and $\mathbf{Q} \in \mathcal{S}$*

$$\mathcal{B}_{\mathcal{R}}(\mathbf{P}, \mathbf{Q}) \geq \mathcal{B}_{\mathcal{R}}(\mathbf{P}, \mathcal{P}_{\Omega}(\mathbf{Q})) + \mathcal{B}_{\mathcal{R}}(\mathcal{P}_{\Omega}(\mathbf{Q}), \mathbf{Q}), \quad (1.11)$$

and in particular

$$\mathcal{B}_{\mathcal{R}}(\mathbf{P}, \mathbf{Q}) \geq \mathcal{B}_{\mathcal{R}}(\mathbf{P}, \mathcal{P}_{\Omega}(\mathbf{Q})). \quad (1.12)$$

We refer the reader to [20], or [22] for a proof of this theorem and further discussions.

1.2.2 Euclidean Projections onto the ℓ_0 and the ℓ_1 -ball

Here, we describe two of key actors in sparse approximation.

Projections onto combinatorial sets: The Euclidean projection of a signal $\mathbf{w} \in \mathbb{R}^N$ on the subspace defined by $\Delta_{\ell_0}(k)$ is provided by:

$$\mathcal{P}_{\Delta_{\ell_0}(k)}(\mathbf{w}) = \underset{\boldsymbol{\alpha} : \boldsymbol{\alpha} \in \Delta_{\ell_0}(k)}{\operatorname{argmin}} \|\boldsymbol{\alpha} - \mathbf{w}\|_2, \quad (1.13)$$

whose solution is hard thresholding. That is, we sort the coefficients of \mathbf{w} in decreasing magnitude and keep the top k and threshold the rest away. This operation can be done in $O(n \log n)$ time complexity via simple sorting routines.

Projections onto convex norms: Given $\mathbf{w} \in \mathbb{R}^N$, the Euclidean projection onto a convex ℓ_1 -norm ball of radius at most τ defines the optimization problem:

$$\mathcal{P}_{\Delta_{\ell_1}(\tau)}(\mathbf{w}) = \underset{\boldsymbol{\alpha} : \boldsymbol{\alpha} \in \Delta_{\ell_1}(\tau)}{\operatorname{argmin}} \|\boldsymbol{\alpha} - \mathbf{w}\|_2, \quad (1.14)$$

whose solution is soft thresholding. That is, we decrease the magnitude of all the coefficients by a constant value just enough to meet the ℓ_1 norm constraint. A solution can be obtained in $O(n \log n)$ time complexity with simple sorting routines, similar to above.

1.2.3 Restricted Isometry Property

In order to establish stronger theoretical guarantees for the algorithms, it is necessary to use Restricted Isometry Property (RIP) assumption. For each positive integers q and k , and each ϵ in $(0, 1)$, an $M \times N$ matrix Φ satisfies the (k, ϵ) RIP in ℓ_q norm $((k, \epsilon)$ RIP- q) [23, 24], if for every k -sparse vector α ,

$$(1 - \epsilon)\|\alpha\|_q \leq \|\Phi\alpha\|_q \leq (1 + \epsilon)\|\alpha\|_q.$$

This assumption implies near isometric embedding of the sparse vectors by the matrix Φ . We just briefly mention that such matrices can be constructed *randomly* using certain classes of distributions [24].

1.3 The GAME Algorithm

1.3.1 A Game Theoretic Reformulation of Sparse Approximation

We start by defining a zero-sum game and then proving that the sparse approximation problem of Equation (1.8) can be reformulated as a zero-sum game.

Definition 1.5 (Zero-sum games [25]). *Let \mathcal{A} and \mathcal{B} be two closed sets. Let $\mathcal{L} : \mathcal{A} \times \mathcal{B} \rightarrow \mathbb{R}$ be a function. The value of a zero sum game, with domains \mathcal{A} and \mathcal{B} with respect to a function \mathcal{L} is defined as*

$$\min_{\mathbf{a} \in \mathcal{A}} \max_{\mathbf{b} \in \mathcal{B}} \mathcal{L}(\mathbf{a}, \mathbf{b}). \quad (1.15)$$

The function \mathcal{L} is usually called the *loss function*. A zero-sum game can be viewed as a game between two players Mindy and Max in the following way. First, Mindy finds a vector \mathbf{a} , and then Max finds a vector \mathbf{b} . The loss that Mindy suffers² is $\mathcal{L}(\mathbf{a}, \mathbf{b})$. The game-value of a zero-sum game is then the loss that Mindy suffers if both Mindy and Max play with their optimal strategies.

Von Neumann's well-known Minimax Theorem [26, 27] states that if both \mathcal{A} and \mathcal{B} are convex compact sets, and if the loss function $\mathcal{L}(\mathbf{a}, \mathbf{b})$ is convex with respect to \mathbf{a} , and concave with respect to \mathbf{b} , then the game-value is independent of the ordering of the game players.

Theorem 1.6 (Von Neumann's Minimax Theorem [26]). *Let \mathcal{A} and \mathcal{B} be closed*

2. which is equal to the gain that Max obtains as the game is zero-sum.

convex sets, and let $\mathcal{L} : \mathcal{A} \times \mathcal{B} \rightarrow \mathbb{R}$ be a function which is convex with respect to its first argument, and concave with respect to its second argument. Then

$$\inf_{\mathbf{a} \in \mathcal{A}} \sup_{\mathbf{b} \in \mathcal{B}} \mathcal{L}(\mathbf{a}, \mathbf{b}) = \sup_{\mathbf{b} \in \mathcal{B}} \inf_{\mathbf{a} \in \mathcal{A}} \mathcal{L}(\mathbf{a}, \mathbf{b}).$$

For the history of the Minimax Theorem see [28]. The Minimax Theorem tells us that for a large class of functions \mathcal{L} , the values of the min-max game in which Mindy goes first is identical to the value of the max-min game in which Max starts the game. The proof of the Minimax Theorem is provided in [29].

Having defined a zero-sum game, and the Von Neumann Minimax Theorem, we next show how the sparse approximation problem of Equation (1.8) can be reformulated as a zero-sum game. Let $p \doteq \frac{q}{q-1}$, and define

$$\Xi_p \doteq \{\mathbf{P} \in \mathbb{R}^M : \|\mathbf{P}\|_p \leq 1\}. \quad (1.16)$$

Define the loss function $\mathcal{L} : \Xi_p \times \Delta_{\ell_1}(\tau) \rightarrow \mathbb{R}$ as

$$\mathcal{L}(\mathbf{P}, \boldsymbol{\alpha}) \doteq \langle \mathbf{P}, (\Phi \boldsymbol{\alpha} - \mathbf{f}) \rangle. \quad (1.17)$$

Observe that the loss-function is bilinear. Now it follows from Hölder inequality that for every $\boldsymbol{\alpha}$ in $\Delta_{\ell_0, \ell_1}(k, \tau)$, and for every \mathbf{P} in Ξ_p

$$\mathcal{L}(\mathbf{P}, \boldsymbol{\alpha}) = \langle \mathbf{P}, (\Phi \boldsymbol{\alpha} - \mathbf{f}) \rangle \leq \|\mathbf{P}\|_p \|\Phi \boldsymbol{\alpha} - \mathbf{f}\|_q \leq \|\Phi \boldsymbol{\alpha} - \mathbf{f}\|_q. \quad (1.18)$$

The inequality of Equation (1.18) becomes equality for

$$P_i^* = \frac{(\Phi \boldsymbol{\alpha} - \mathbf{f})_i^{q/p}}{\left(\sum_{i=1}^M (\Phi \boldsymbol{\alpha} - \mathbf{f})_i^q \right)^{1/p}}.$$

Therefore

$$\max_{\mathbf{P} \in \Xi_p} \mathcal{L}(\mathbf{P}, \boldsymbol{\alpha}) = \max_{\mathbf{P} \in \Xi_p} \langle \mathbf{P}, (\Phi \boldsymbol{\alpha} - \mathbf{f}) \rangle = \langle \mathbf{P}^*, (\Phi \boldsymbol{\alpha} - \mathbf{f}) \rangle = \|\Phi \boldsymbol{\alpha} - \mathbf{f}\|_q. \quad (1.19)$$

Equation (1.19) is true for every $\boldsymbol{\alpha} \in \Delta_{\ell_1}(\tau)$. As a result, by taking the minimum over $\Delta_{\ell_0, \ell_1}(k, \tau)$ we get

$$\min_{\boldsymbol{\alpha} \in \Delta_{\ell_0, \ell_1}(k, \tau)} \|\Phi \boldsymbol{\alpha} - \mathbf{f}\|_q = \min_{\boldsymbol{\alpha} \in \Delta_{\ell_0, \ell_1}(k, \tau)} \max_{\mathbf{P} \in \Xi_p} \mathcal{L}(\mathbf{P}, \boldsymbol{\alpha}).$$

Similarly by taking the minimum over $\Delta_{\ell_1}(\tau)$ we get

$$\min_{\boldsymbol{\alpha} \in \Delta_{\ell_1}(\tau)} \|\Phi \boldsymbol{\alpha} - \mathbf{f}\|_q = \min_{\boldsymbol{\alpha} \in \Delta_{\ell_1}(\tau)} \max_{\mathbf{P} \in \Xi_p} \mathcal{L}(\mathbf{P}, \boldsymbol{\alpha}). \quad (1.20)$$

Solving the sparse approximation problem of Equation (1.8) is therefore equivalent to finding the optimal strategies of the game

$$\min_{\boldsymbol{\alpha} \in \Delta_{\ell_0, \ell_1}(k, \tau)} \max_{\mathbf{P} \in \Xi_p} \mathcal{L}(\mathbf{P}, \boldsymbol{\alpha}). \quad (1.21)$$

In the next section we provide a primal-dual algorithm that approximately solves this min-max game. Observe that since $\Delta_{\ell_0, \ell_1}(k, \tau)$ is a subset of $\Delta_{\ell_1}(\tau)$, we always have

$$\min_{\alpha \in \Delta_{\ell_1}(\tau)} \max_{\mathbf{P} \in \Xi_p} \mathcal{L}(\mathbf{P}, \alpha) \leq \min_{\alpha \in \Delta_{\ell_0, \ell_1}(k, \tau)} \max_{\mathbf{P} \in \Xi_p} \mathcal{L}(\mathbf{P}, \alpha),$$

and therefore, in order to approximately solve the game of Equation (1.21), it is sufficient to find $\hat{\alpha} \in \Delta_{\ell_0, \ell_1}(k, \tau)$ with

$$\max_{\mathbf{P} \in \Xi_p} \mathcal{L}(\mathbf{P}, \hat{\alpha}) \approx \min_{\alpha \in \Delta_{\ell_1}(\tau)} \max_{\mathbf{P} \in \Xi_p} \mathcal{L}(\mathbf{P}, \alpha). \quad (1.22)$$

1.3.2 Algorithm Description

In this section we provide an efficient algorithm for approximately solving the problem of sparse approximation in ℓ_q norm, defined by Equation (1.10). Let $\mathcal{L}(\mathbf{P}, \alpha)$ be the loss function defined by Equation (1.17), and recall that in order to approximately solve Equation (1.10), it is sufficient to find a sparse vector $\hat{\alpha} \in \Delta_{\ell_0, \ell_1}(k, \tau)$ such that

$$\max_{\mathbf{P} \in \Xi_p} \mathcal{L}(\mathbf{P}, \hat{\alpha}) \approx \min_{\alpha' \in \Delta_{\ell_1}(\tau)} \max_{\mathbf{P} \in \Xi_p} \mathcal{L}(\mathbf{P}, \alpha'). \quad (1.23)$$

The original sparse approximation problem of Equation (1.10) is NP-complete, but it is computationally feasible to compute the value of the min-max game

$$\min_{\alpha' \in \Delta_{\ell_1}(\tau)} \max_{\mathbf{P} \in \Xi_p} \mathcal{L}(\mathbf{P}, \alpha'). \quad (1.24)$$

The reason is that the loss function $\mathcal{L}(\mathbf{P}, \alpha)$ of Equation (1.17) is a bilinear function, and the sets $\Delta_{\ell_1}(\tau)$, and Ξ_p are both convex and closed.

Therefore, finding the game values and optimal strategies of the game of Equation (1.24) is equivalent to solving a convex optimization problem and can be done using off-the-shelf non-smooth convex optimization methods [30, 31]. However, if an off-the-shelf convex optimization method is used, then there is no guarantee that the recovered strategy $\hat{\alpha}$ is also sparse. We need an approximation algorithm that finds near-optimal strategies $\hat{\alpha}$ and $\hat{\mathbf{P}}$ for Mindy and Max with the additional guarantee that Mindy's near optimal strategy $\hat{\alpha}$ is sparse.

Here we introduce the *Game-theoretic Approximate Matching Estimator* (GAME) algorithm which finds a sparse *approximation* to the min-max optimal solution of the game defined in Equation (1.24). The GAME algorithm relies on the general primal-dual approach which was originally applied to developing strategies for repeated games [29] (see also [32] and [33]). The pseudocode of the GAME Algorithm is provided in Algorithm 1.1.

The GAME Algorithm can be viewed as a repeated game between two players Mindy and Max who iteratively update their current strategies \mathbf{P}^t and α^t , with the aim of ultimately finding near-optimal strategies based on a T -round interaction

Algorithm 1.1 GAME Algorithm for Sparse Approximation in ℓ_q -norm.

Inputs: M -dimensional vector \mathbf{f} , $M \times N$ matrix Φ , number of iterations T , sparse approximation norm q , Bregman function \mathcal{R} and regularization parameter η .

Output: N -dimensional vector $\hat{\alpha}$

with each other. Here, we briefly explain how each player updates his/her current strategy based on the new update from the other player.

Recall that the ultimate goal is to find the solution of the game

$$\min_{\alpha' \in \Delta_{\ell_1}(\tau)} \max_{\mathbf{P} \in \Xi_p} \mathcal{L}(\mathbf{P}, \alpha').$$

At the begining of each iteration t , Mindy receives the updated value \mathbf{P}^t from Max. A greedy Mindy only focuses on Max's current strategy, and updates her current strategy to $\alpha^t = \arg \min_{\alpha \in \Delta_{\ell_1}(\tau)} \mathcal{L}(\mathbf{P}^t, \alpha)$. In the following lemma we show that this is indeed what our Mindy does in the first three steps of the main loop.

Lemma 1.7. *Let \mathbf{P}^t denote Max's strategy at the begining of iteration t . Let $\mathbf{r}^t = \Phi^\top \mathbf{P}^t$, and let i denote the index of a largest (in magnitude) element of \mathbf{r}^t . Let α^t be a 1-sparse vector with $\text{Supp}(\alpha^t) = \{i\}$ and with $\alpha_i^t = -\tau \text{Sign}(r_i^t)$. Then $\alpha^t = \arg \min_{\alpha \in \Delta_{\ell_1}(\tau)} \mathcal{L}(\mathbf{P}^t, \alpha)$.*

Proof. Let $\tilde{\alpha}$ be any solution $\tilde{\alpha} = \arg \min_{\alpha \in \Delta_{\ell_1}(\tau)} \mathcal{L}(\mathbf{P}^t, \alpha)$. It follows from the bilinearity of the loss function (Equation (1.17)) that

$$\begin{aligned} \tilde{\alpha} &= \arg \min_{\alpha \in \Delta_{\ell_1}(\tau)} \mathcal{L}(\mathbf{P}^t, \alpha) \\ &= \arg \min_{\alpha \in \Delta_{\ell_1}(\tau)} \langle \mathbf{P}^t, \Phi \alpha - \mathbf{f} \rangle = \arg \min_{\alpha \in \Delta_{\ell_1}(\tau)} \langle \Phi^\top \mathbf{P}^t, \alpha \rangle. \end{aligned}$$

Hence, Hölder inequality yields that for every $\alpha^\# \in \Delta_{\ell_1}(\tau)$,

$$\langle \Phi^\top \mathbf{P}^t, \alpha^\# \rangle \geq -\|\alpha^\#\|_1 \|\Phi^\top \mathbf{P}^t\|_\infty \geq -\tau \|\Phi^\top \mathbf{P}^t\|_\infty. \quad (1.25)$$

Now let α^t be a 1-sparse vector with $\text{Supp}(\alpha^t) = \{i\}$ and $\alpha_i^t = -\tau \text{Sign}(r_i^t)$. Then $\alpha^t \in \Delta_{\ell_1}(\tau)$, and

$$\langle \Phi^\top \mathbf{P}^t, \alpha^t \rangle = -\tau \|\Phi^\top \mathbf{P}^t\|_\infty.$$

In other words, for α^t the Holder inequality is an equality. Hence α^t is a minimizer of $\langle \Phi^\top \mathbf{P}^t, \alpha \rangle$. \square

Thus far we have seen that at each iteration Mindy always finds a 1-sparse solution $\alpha^t = \arg \min_{\alpha \in \Delta_{\ell_1}(\tau)} \mathcal{L}(\mathbf{P}^t, \alpha)$. Mindy then sends her updated strategy α^t to Max, and now it is Max's turn to update his strategy. A greedy Max would prefer to update his strategy as $\mathbf{P}^{t+1} = \arg \max_{\mathbf{P} \in \Xi_p} \mathcal{L}(\mathbf{P}, \alpha^t)$. However, our Max is more conservative and prefers to stay close to his previous value \mathbf{P}^t . In other words, Max has two competing objectives

1. Maximizing $\mathcal{L}(\mathbf{P}, \boldsymbol{\alpha}^t)$, or equivalently minimizing $-\mathcal{L}(\mathbf{P}, \boldsymbol{\alpha}^t)$.
2. Remaining close to the previous strategy \mathbf{P}^t , by minimizing $\mathcal{B}_{\mathcal{R}}(\mathbf{P}, \mathbf{P}^t)$.

Let

$$\mathcal{L}_{\mathcal{R}}(\mathbf{P}) \doteq -\eta\mathcal{L}(\mathbf{P}, \boldsymbol{\alpha}^t) + \mathcal{B}_{\mathcal{R}}(\mathbf{P}, \mathbf{P}^t),$$

be a regularized loss function which is a linear combination of the two objectives above.

A conservative Max then tries to minimize a combination of the two objectives above by minimizing the regularized loss function

$$\mathbf{P}^{t+1} = \arg \min_{\mathbf{P} \in \Xi_p} \mathcal{L}_{\mathcal{R}}(\mathbf{P}) = \arg \min_{\mathbf{P} \in \Xi_p} -\eta\mathcal{L}(\mathbf{P}, \boldsymbol{\alpha}^t) + \mathcal{B}_{\mathcal{R}}(\mathbf{P}, \mathbf{P}^t). \quad (1.26)$$

Unfortunately, it is not so easy to efficiently solve the optimization problem of Equation (1.26) at every iteration. To overcome this difficulty, our Max first ignores the constraint $\mathbf{P}^{t+1} \in \Xi_p$, and instead finds a global optimizer of $\mathcal{L}_{\mathcal{R}}(\mathbf{P})$ by setting $\nabla \mathcal{L}_{\mathcal{R}}(\mathbf{P}) = \mathbf{0}_M$, and then projects back the result to Ξ_p via a Bregman projection.

More precisely, it follows from the Property (P2) of Bregman distance (Theorem 1.2) that for every \mathbf{P}

$$\nabla \mathcal{L}_{\mathcal{R}}(\mathbf{P}) = -\eta(\Phi \boldsymbol{\alpha}^t - \mathbf{f}) + \nabla \mathcal{R}(\mathbf{P}) - \nabla \mathcal{R}(\mathbf{P}^t),$$

and therefore if \mathbf{Q}^t is a point with

$$\nabla \mathcal{R}(\mathbf{Q}^t) = \nabla \mathcal{R}(\mathbf{P}^{t-1}) + \eta(\Phi \boldsymbol{\alpha}^t - \mathbf{f}),$$

then $\nabla \mathcal{L}_{\mathcal{R}}(\mathbf{Q}^t) = \mathbf{0}_M$.

The vector \mathbf{Q}^t is finally projected back to Ξ_p via a Bregman projection to ensure that Max's new strategy is in the feasible set Ξ_p .

1.3.3 The GAME Guarantees

In this section we prove that the GAME algorithm finds a near-optimal solution for the sparse approximation problem of Equation (1.10). The analysis of the GAME algorithm relies heavily on the analysis of the generic primal-dual approach. This approach originates from the *link-function methodology* in computational optimization [33, 34], and is related to the *mirror descent* approach in the optimization community [35, 36]. The primal-dual Bregman optimization approach is widely used in online optimization applications including portfolio selection [37, 38], online learning [39], and boosting [40, 41].

However, there is a major difference between the sparse approximation problem and the problem of online convex optimization. In the sparse approximation problem, the set $\mathcal{A} = \Delta_{\ell_0, \ell_1}(k, \tau)$ is not convex anymore; therefore, there is no guarantee that an online convex optimization algorithm outputs a sparse strategy $\hat{\boldsymbol{\alpha}}$. Hence, it is not possible to directly translate the bounds from the online convex optimization

scheme to the sparse approximation scheme.

Moreover, as discussed in Lemma 1.7, there is also a major difference between the Mindy players of the GAME algorithm and the general Mindy of general online convex optimization games. In the GAME algorithm, Mindy is not a *blackbox adversary* that responds with an update to her strategy based on Max's update. Here, Mindy always performs a greedy update and finds the *best* strategy as a response to Max's update. Moreover, our Mindy always finds a 1-sparse new strategy. That is, she looks among all best responses to Max's update, and finds a 1-sparse strategy among them.

As we will see next, the combination of cooperativeness by Mindy, and standard ideas for bounding the regret in online convex optimization schemes, enables us to analyze the GAME algorithm for sparse approximation. The following lemma bounds the regret loss of the primal-dual strategy in online convex optimization problems and is proved in [32].

Theorem 1.8. *Let q and T be positive integers, and let $p = \frac{q}{q-1}$. Suppose that \mathcal{R} is such that for every $\mathbf{P}, \mathbf{Q} \in \Xi_p$, $\mathcal{B}_{\mathcal{R}}(\mathbf{P}, \mathbf{Q}) \geq \|\mathbf{P} - \mathbf{Q}\|_p^2$, and let*

$$G = \max_{\boldsymbol{\alpha} \in \Delta_{\ell_0, \ell_1}(1, \tau)} \|\Phi \boldsymbol{\alpha} - \mathbf{f}\|_q. \quad (1.27)$$

Also assume that for every $\mathbf{P} \in \Xi_p$, we have $\mathcal{B}_{\mathcal{R}}(\mathbf{P}, \mathbf{P}^1) \leq D^2$. Suppose

$$\langle (\mathbf{P}^1, \boldsymbol{\alpha}^1), \dots, (\mathbf{P}^T, \boldsymbol{\alpha}^T) \rangle$$

is the sequence of pairs generated by the GAME Algorithm after T iterations with $\eta = \frac{2D}{G\sqrt{T}}$. Then

$$\max_{\mathbf{P} \in \Xi_p} \frac{1}{T} \sum_{t=1}^T \mathcal{L}(\mathbf{P}, \boldsymbol{\alpha}^t) \leq \frac{1}{T} \sum_{t=1}^T \mathcal{L}(\mathbf{P}^t, \boldsymbol{\alpha}^t) + \frac{DG}{2\sqrt{T}}.$$

Proof. The proof of Theorem 1.8 is based on the geometric properties of the Bregman functions, and is provided in [32]. \square

Next we use Theorem 1.8 to show that the GAME algorithm after T iterations finds a T -sparse vector $\hat{\boldsymbol{\alpha}}$ with near-optimal value $\|\Phi \hat{\boldsymbol{\alpha}} - \mathbf{f}\|_q$.

Theorem 1.9. *Let q and T be positive integers, and let $p = \frac{q}{q-1}$. Suppose that for every $\mathbf{P}, \mathbf{Q} \in \Xi_p$, the function \mathcal{R} satisfies $\mathcal{B}_{\mathcal{R}}(\mathbf{P}, \mathbf{Q}) \geq \|\mathbf{P} - \mathbf{Q}\|_p^2$, and let*

$$G = \max_{\boldsymbol{\alpha} \in \Delta_{\ell_0, \ell_1}(1, \tau)} \|\Phi \boldsymbol{\alpha} - \mathbf{f}\|_q. \quad (1.28)$$

Also assume that for every $\mathbf{P} \in \Xi_p$, we have $\mathcal{B}_{\mathcal{R}}(\mathbf{P}, \mathbf{P}^1) \leq D^2$. Suppose

$$\langle (\mathbf{P}^1, \boldsymbol{\alpha}^1), \dots, (\mathbf{P}^T, \boldsymbol{\alpha}^T) \rangle$$

is the sequence of pairs generated by the GAME Algorithm after T iterations with $\eta = \frac{2D}{G\sqrt{T}}$. Let $\hat{\boldsymbol{\alpha}} = \frac{1}{T} \sum_{t=1}^T \boldsymbol{\alpha}^t$ be the output of the GAME algorithm. Then $\hat{\boldsymbol{\alpha}}$ is a

T -sparse vector with $\|\hat{\alpha}\|_1 \leq \tau$ and

$$\|\Phi \hat{\alpha} - \mathbf{f}\|_q \leq \min_{\alpha \in \Delta_{\ell_0, \ell_1}(T, \tau)} \|\Phi \alpha - \mathbf{f}\|_q + \frac{\text{DG}}{2\sqrt{T}}. \quad (1.29)$$

Proof. It follows from Step 2. of Algorithm 1.1 that every α^t is 1-sparse and $\|\alpha^t\|_1 = \tau$. Therefore, $\hat{\alpha} = \frac{1}{T} \sum_{t=1}^T \alpha^t$ can have at most T non-zero entries and moreover $\|\hat{\alpha}\|_1 \leq \frac{1}{T} \sum_{t=1}^T \|\alpha^t\|_1 \leq \tau$. Therefore $\hat{\alpha}$ is in $\Delta_{\ell_0, \ell_1}(T, \tau)$.

Next we show that the Equation 1.29 holds for $\hat{\alpha}$. Let $\hat{\mathbf{P}} = \frac{1}{T} \sum_{t=1}^T \mathbf{P}^t$. Observe that

$$\begin{aligned} \min_{\alpha \in \Delta_{\ell_1}(\tau)} \max_{\mathbf{P} \in \Xi_p} \mathcal{L}(\mathbf{P}, \alpha) &\stackrel{(e)}{=} \max_{\mathbf{P} \in \Xi_p} \min_{\alpha \in \Delta_{\ell_1}(\tau)} \mathcal{L}(\mathbf{P}, \alpha) \\ &\stackrel{(f)}{\geq} \min_{\alpha \in \Delta_{\ell_1}(\tau)} \mathcal{L}(\hat{\mathbf{P}}, \alpha) \\ &\stackrel{(g)}{\geq} \frac{1}{T} \min_{\alpha \in \Delta_{\ell_1}(\tau)} \sum_{t=1}^T \mathcal{L}(\mathbf{P}^t, \alpha) \\ &\stackrel{(h)}{\geq} \frac{1}{T} \sum_{t=1}^T \min_{\alpha \in \Delta_{\ell_1}(\tau)} \mathcal{L}(\mathbf{P}^t, \alpha) \stackrel{(i)}{=} \frac{1}{T} \sum_{t=1}^T \mathcal{L}(\mathbf{P}^t, \alpha^t) \\ &\stackrel{(j)}{\geq} \max_{\mathbf{P} \in \Xi_p} \mathcal{L}\left(\mathbf{P}, \frac{1}{T} \sum_{t=1}^T \alpha^t\right) - \frac{\text{DG}}{2\sqrt{T}}. \end{aligned}$$

Equality (e) is the minimax Theorem (Theorem 1.6). Inequality (f) follows from the definition of the max function. Inequalities (g) and (h) are consequences of the bilinearity of \mathcal{L} and concavity of the min function. Equality (i) is valid by the definition of α^t , and Inequality (j) follows from Theorem 1.8. As a result

$$\begin{aligned} \|\Phi \hat{\alpha} - \mathbf{f}\|_q &= \max_{\mathbf{P} \in \Xi_p} \mathcal{L}(\mathbf{P}, \hat{\alpha}) \leq \min_{\alpha \in \Delta_{\ell_1}(\tau)} \max_{\mathbf{P} \in \Xi_p} \mathcal{L}(\mathbf{P}, \alpha) \\ &\quad + \frac{\text{DG}}{2\sqrt{T}} = \min_{\alpha \in \Delta_{\ell_1}(\tau)} \|\Phi \alpha - \mathbf{f}\|_q + \frac{\text{DG}}{2\sqrt{T}}. \end{aligned}$$

□

Remark 1.10. In general, different choices for the Bregman function may lead to different convergence bounds with different running times to perform the new projections and updates. For instance, a multiplicative update version of the algorithm can be derived by using the Bregman divergence based on the Kullback-Leibler function, and an additive update version of the algorithm can be derived by using the Bregman divergence based on the squared Euclidean function.

Theorem 1.9 is applicable to *any* sensing matrix. Nevertheless, it does not guarantee that the estimate vector $\hat{\alpha}$ is close enough to the target vector α^* . However, if the sensing matrix satisfies the RIP- q property, then it is possible to bound the data-domain error $\|\hat{\alpha} - \alpha^*\|_q$ as well.

Theorem 1.11. *Let q, k , and T be positive integers, let ϵ be a number in $(0, 1)$, and let $p = \frac{q}{q-1}$. Suppose that for every $\mathbf{P}, \mathbf{Q} \in \Xi_p$, the function \mathcal{R} satisfies $\mathcal{B}_{\mathcal{R}}(\mathbf{P}, \mathbf{Q}) \geq \|\mathbf{P} - \mathbf{Q}\|_p^2$, and let Φ be an $M \times N$ sensing matrix satisfying the $(k + T, \epsilon)$ RIP- q property. Let α^* be a k -sparse vector with $\|\alpha^*\|_1 \leq \tau$, let \mathbf{e}_M be an arbitrary noise vector in \mathbb{R}^M , and set $\mathbf{f} = \Phi\alpha^* + \mathbf{e}_M$. Let G, D , and η be as of Theorem 1.9, and let $\hat{\alpha}$ be the output of the GAME algorithm after T iterations. Then $\hat{\alpha}$ is a T -sparse vector with $\|\hat{\alpha}\|_1 \leq \tau$ and*

$$\|\hat{\alpha} - \alpha^*\|_q \leq \frac{2\|\mathbf{e}_M\|_q + \frac{DG}{2\sqrt{T}}}{(1 - \epsilon)}. \quad (1.30)$$

Proof. Since $\hat{\alpha}$ is T -sparse and α^* is k -sparse, $\hat{\alpha} - \alpha^*$ is $(T + k)$ -sparse. Therefore, it follows from the RIP- q property of the sensing matrix that

$$\begin{aligned} (1 - \epsilon)\|\hat{\alpha} - \alpha^*\|_q &\leq \|\Phi(\hat{\alpha} - \alpha^*)\|_q \leq \|\Phi\hat{\alpha} - \mathbf{f}\|_q + \|\mathbf{e}_M\|_q \\ &\leq \|\Phi\alpha^* - \mathbf{f}\|_q + \frac{DG}{2\sqrt{T}} + \|\mathbf{e}_M\|_q = 2\|\mathbf{e}_M\|_q + \frac{DG}{2\sqrt{T}}. \end{aligned} \quad (1.31)$$

□

1.4 The CLASH Algorithm

1.4.1 Hard Thresholding Formulations of Sparse Approximation

As already stated, solving (1.2) is NP-hard and exhaustive search over $\binom{N}{k}$ possible support set configurations of the k -sparse solution is mandatory. Contrary to this brute-force approach, hard thresholding algorithms [10, 11, 12, 13, 14] navigate through the low-dimensional k -sparse subspaces, *pursuing* an appropriate support set such to minimize the data error in (1.4). To achieve this, these approaches apply greedy support set selection rules to iteratively compute and refine a putative solution α_i using only first-order information $\nabla f(\alpha_{i-1})$ at each iteration i .

Subspace Pursuit (SP) [11] algorithm is a combinatorial greedy algorithm that borrows both from Orthogonal Matching Pursuit (OMP) and Iterative Hard Thresholding [13] (IHT) methods. A sketch of the algorithm is given in Algorithm 2. The basic idea behind SP consists in looking for a good support set by iteratively collecting an extended candidate support set $\hat{\mathcal{A}}_i$ with $|\hat{\mathcal{A}}_i| \leq 2k$ (Step 4) and then finding the k -sparse vector α_{i+1} that best fits the measurements within the restricted support set $\hat{\mathcal{A}}_i$, i.e., the support set α_{i+1} satisfies $\mathcal{A}_{i+1} \triangleq \text{supp}(\alpha_{i+1}) \subseteq \hat{\mathcal{A}}_i$ (Steps 5-6).

Algorithm 1.2 Subspace Pursuit Algorithm

Input: $\mathbf{f}, \Phi, k, \text{MaxIter}$. **Output:** $\hat{\alpha} \leftarrow \arg\min_{\mathbf{v}: \text{supp}(\mathbf{v}) \subseteq \mathcal{A}_i} \|\mathbf{f} - \Phi\mathbf{v}\|_2^2$

In [42], Foucart improves the initial RIP conditions of SP algorithm, which we present here as a corollary:

Corollary 1.12 (SP Iteration Invariant). *SP algorithm satisfies the following recursive formula:*

$$\|\alpha_{i+1} - \alpha^*\|_2 \leq \rho \|\alpha_i - \alpha^*\|_2 + c \|\mathbf{n}\|_2, \quad (1.32)$$

where $c = \sqrt{\frac{2(1+3\delta_{3k}^2)}{1-\delta_{3k}}} + \frac{\sqrt{(1+3\delta_{3k}^2)(1+\delta_{2k})}}{1-\delta_{3k}} + \sqrt{3(1+\delta_{2k})}$ and $\rho < 1$ given that $\delta_{3k} < 0.38427$.

1.4.2 Algorithm Description

In this section, we expose CLASH algorithm, a Subspace Pursuit [11] variant, as a running example for our subsequent developments. We underline that norm constraints can be also incorporated into alternative state-of-the-art hard thresholding frameworks [10, 11, 12, 13, 14].

Algorithm 1.3 The CLASH Algorithm

Input: $f, \Phi, \Delta_{\ell_0, \ell_1}(k, \tau)$, Tolerance, MaxIterations **Output:** α_i .

The CLASH algorithm approximates α^* according to the optimization formulation (1.8) where $q = 2$. We provide a pseudo-code of an example implementation of CLASH in Algorithm 1.3. To complete the i -th iteration, CLASH initially identifies a $2k$ extended support set $\hat{\mathcal{A}}_i$ to explore via the *Active set expansion* step (Step 1)—the set $\hat{\mathcal{A}}_i$ is constituted by the union of the support \mathcal{A}_i of the current solution α_i and an additional k -sparse support where the projected gradient onto $\Delta_{\ell_0}(k)$ can make most impact on the loading vector, *complementary* to \mathcal{A}_i . Given $\hat{\mathcal{A}}_i$, the *Greedy descent with least absolute shrinkage* step (Step 2) solves a least-squares problem over ℓ_1 -norm constraint to decrease the data error $f(\alpha)$, restricted over the active support set $\hat{\mathcal{A}}_i$. In sequence, we project the $2k$ -sparse solution of Step 2 onto $\Delta_{\ell_0}(k)$ to arbitrate the active support set via the *Combinatorial selection* step (Step 3). Finally, CLASH de-biases the result on the putative solution support using the *De-bias* step (Step 4).

1.4.3 The CLASH Guarantees

CLASH iterations satisfy the following worst-case guarantee:

Theorem 1.13. *[Iteration invariant] Let α^* be the true solution. Then, the i -th iterate α_i of CLASH satisfies the following recursion*

$$\|\alpha_{i+1} - \alpha^*\|_2 \leq \rho \|\alpha_i - \alpha^*\|_2 + c_1(\delta_{2k}, \delta_{3k}) \|\mathbf{n}\|_2, \text{ where} \quad (1.33)$$

$$c_1(\delta_{2k}, \delta_{3k}) \triangleq \frac{1}{\sqrt{1 - \delta_{2k}^2}} \left(\sqrt{1 + 3\delta_{3k}^2} \left(\sqrt{\frac{2(1 + \delta_{3k})}{1 - \delta_{3k}^2}} + \frac{\sqrt{1 + \delta_{2k}}}{1 - \delta_{3k}} \right) + \sqrt{3(1 + \delta_{2k})} \right) + \frac{\sqrt{1 + \delta_k}}{1 - \delta_{2k}}, \quad (1.34)$$

and $\rho \triangleq \frac{\delta_{3k} + \delta_{2k}}{\sqrt{1 - \delta_{2k}^2}} \sqrt{\frac{1 + 3\delta_{3k}^2}{1 - \delta_{3k}^2}}$. Moreover, when $\delta_{3k} < 0.3658$, the iterations are contractive (i.e., $\rho < 1$).

A detailed proof of Theorem 1.13 can be found in [19]. Theorem 1.13 shows that the isometry requirements of CLASH are competitive with those of mainstream hard thresholding methods, such as SP, even though CLASH incorporates the ℓ_1 -norm constraints—furthermore, we observe improved signal reconstruction performance compared to these methods, as shown in the Experiments section.

1.5 Experiments

In this section, we provide experimental results to demonstrate the performances of the GAME and CLASH Algorithms.

1.5.1 Performance of the ℓ_∞ GAME algorithm

In this experiment, we fix $N = 1000$, $M = 200$ and $k = 20$, and generate a 200×1000 Gaussian matrix Φ . Each experiment is repeated independently 50 times. We compare the performance of the ℓ_∞ GAME algorithm, which approximately solves the non-convex problem

$$\underset{\alpha \in \Delta_{\ell_0, \ell_1}(k, \tau)}{\text{minimize}} \quad \|\Phi^\top \Phi \alpha - \Phi^\top \mathbf{f}\|_\infty \quad (1.35)$$

with state-of-the-art Dantzig Selector solvers [43, 44] that solve linear optimization

$$\underset{\alpha \in \Delta_{\ell_1}(\tau)}{\text{minimize}} \quad \|\Phi^\top \Phi \alpha - \Phi^\top \mathbf{f}\|_\infty \quad (1.36)$$

The compressive measurements were generated in the presence of white Gaussian noise. The noise vector consists of M iid $\mathcal{N}(0, \sigma^2)$ elements, where σ ranges from $10^{-3.5}$ to $10^{-0.5}$. Figure 1.3 compares the data-domain ℓ_2 -error ($\|\alpha^* - \hat{\alpha}\|_2 / \|\alpha^*\|_2$) of the GAME algorithm with the error of ℓ_1 -magic algorithm [45] and the Homotopy algorithm [46] which are state-of-the-art Dantzig Selector optimizers. As illustrated in Figure 1.3, as σ increases to 10^{-3} , the GAME algorithm outperforms the ℓ_1 -magic and Homotopy algorithms.




Figure 1.3: Signal approximation experiments with ℓ_1 -magic, Homotopy, and GAME algorithms. The measurement noise standard deviation ranges from $10^{-3.5}$ to $10^{-0.5}$, and the approximation error is measured as $\|\alpha^* - \hat{\alpha}\|_2$.

1.5.2 Performance of Clash Algorithm

Noise resilience: We generate random realizations of the model $\mathbf{f} = \Phi\alpha^*$ for $N = 1000$, $M = 305$ and $k = 115$ where k is known a-priori and α^* admits the simple sparsity model. We construct α^* as a k -sparse vector with iid $\mathcal{N}(0, 1)$ elements with $\|\alpha^*\|_2 = 1$. We repeat the same experiment independently for 50 Monte-Carlo iterations. In this experiment, we examine the signal recovery performance of CLASH compared to the following state-of-the-art methods: *i*) Lasso (1.4) as a projected gradient method, *ii*) Basis Pursuit [15] using SPGL1 implementation [47] and, *iii*) Subspace Pursuit [11]. We test the recovery performance of the aforementioned methods for various noise standard deviations – the empirical results are depicted in Figure 1.4. We observe that the combination of hard thresholding with norm constraints significantly *improves* the signal recovery performance over both convex- and combinatorial-based approaches.

Improved recovery using CLASH: We generate random realizations of the model $\mathbf{f} = \Phi\alpha^* + \mathbf{n}$ for $N = 500$, $M = 160$ and $k = \{57, 62\}$ for the noisy and the noiseless case respectively, where k is known a-priori. We construct α^* as a k -sparse vector with iid $\mathcal{N}(0, 1)$ elements with $\|\alpha^*\|_2 = 1$. In the noisy case, we assume $\|\mathbf{n}\|_2 = 0.05$. We perform 500 independent Monte-Carlo iterations. We then sweep τ and then examine the signal recovery performance of CLASH compared to the same methods above. Note that, if τ is large, norm constraints have no impact in recovery and CLASH must admit identical performance to SP.

Figure 1.5 illustrates that the combination of hard thresholding with norm constraints can *improve* the signal recovery performance significantly over convex-only and hard thresholding-only methods. CLASH perfectly recovers the signal when




Figure 1.4: Signal approximation experiments with CLASH, Lasso, and BP algorithms. The measurement noise standard deviation ranges from 10^{-5} to 10^{-1} , and the approximation error is measured as $\|\alpha^* - \hat{\alpha}\|_2$.




Figure 1.5: Improved signal recovery using CLASH.

the regularization parameter is close to $\|\alpha^*\|_1$. When $\tau \ll \|\alpha^*\|_1$ or $\tau \gg \|\alpha^*\|_1$, the performance degrades.

1.6 Conclusions

We discussed two sparse recovery algorithms that explicitly leverage convex ℓ_1 and non-convex ℓ_0 priors jointly. While the ℓ_1 prior is conventionally motivated as the “convexification” of the ℓ_0 prior, we saw that this interpretation is incomplete: it actually is a convexification of the ℓ_0 -constrained set with a maximum scale. We also discovered that the interplay of these two—seemingly related—priors could lead to not only strong theoretical recovery guarantees from weaker assumptions than commonly used in sparse recovery, but also improved empirical performance over the existing solvers. To obtain our results, we reviewed some important topics from game theory, convex and combinatorial optimization literature. We believe that

understanding and exploiting the interplay of such convex and non-convex priors could lead to radically new, scalable regression approaches, which can leverage decades of work in diverse theoretical disciplines.

Acknowledgements

VC and AK's work was supported in part by the European Commission under Grant MIRG-268398, ERC Future Proof, DARPA KeCoM program # 11-DARPA-1055, and SNF 200021-132548 grants. VC also would like to acknowledge Rice University for his Faculty Fellowship. SJ thanks Robert Calderbank and Rob Schapire for providing insightful comments.

Bibliography

1. E. J. Candès, J. K. Romberg, and T. Tao, "Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information," *IEEE Transactions on Information Theory*, vol. 52, pp. 489–509, 2006.
2. D. L. Donoho, "Compressed sensing," *IEEE Transactions on Information Theory*, vol. 52, no. 4, pp. 1289–1306, 2006.
3. R. Tibshirani, "Regression shrinkage and selection via the LASSO," *J. Royal. Statist. Soc.*, vol. 5, no. 1, pp. 267–288, 1996.
4. M. J. Wainwright, "Sharp thresholds for high-dimensional and noisy sparsity recovery using ℓ_1 -constrained quadratic programming (lasso)," *IEEE Transactions on Information Theory*, vol. 55, no. 5, pp. 2183–2202, 2009.
5. A. J. Miller, *Subset selection in regression*. New York: Chapman-Hall, 1990.
6. P. Ravikumar, M. J. Wainwright, and J. Lafferty, "High-dimensional Ising model selection using ℓ_1 -regularized logistic regression," *Annals of Statistics*, vol. 38, pp. 1287–1319, 2010.
7. N. Meinshausen and P. Bühlmann, "High dimensional graphs and variable selection with the lasso," *Annals of Statistics*, vol. 34, pp. 1436–1462, 2006.
8. D. Paul, "Asymptotics of sample eigenstructure for a large-dimensional spiked covariance model," *Statistica Sinica*, vol. 17, p. 16171642, 2007.
9. M. Johnstone and A. Lu, "On consistency and sparsity for principal components analysis in high dimensions," *Journal of the American Statistical Association*, vol. 104, no. 486, p. 682693, 2009.
10. A. Kyrillidis and V. Cevher, "Recipes for hard thresholding methods," *Technical Report*, 2011.
11. W. Dai and O. Milenkovic, "Subspace pursuit for compressive sensing signal reconstruction," *IEEE Transactions on Information Theory*, vol. 5, 2009.
12. D. Needell and J. Tropp, "CoSaMP: Iterative signal recovery from incomplete and inaccurate samples," *Applied and Computational Harmonic Analysis*, vol. 26, no. 3, pp. 16–42, 2007.

13. T. Blumensath and M. E. Davies, "Iterative hard thresholding for compressed sensing," *Applied and Computational Harmonic Analysis*, vol. 27, no. 3, pp. 265–274, 2009.
14. S. Foucart, "Hard thresholding pursuit: An algorithm for compressive sensing, preprint," 2011.
15. S. S. Chen, D. L. Donoho, and M. A. Saunders, "Atomic Decomposition by Basis Pursuit," *SIAM Journal on Scientific Computing*, vol. 20, p. 33, 1998.
16. R. Tibshirani, "Regression shrinkage and selection via the lasso," *J. Royal. Statist. Soc B*, vol. 58, no. 1, pp. 267–288, 1996.
17. G. P. A. Kyrillidis and V. Cevher, "Hard thresholding with norm constraints," *Technical Report*, 2011.
18. R. G. Baraniuk, V. Cevher, M. F. Duarte, and C. Hegde, "Model-based compressive sensing," *IEEE Transactions on Information Theory*, vol. 56, 2010.
19. A. Kyrillidis and V. Cevher, "Combinatorial selection and least absolute shrinkage via the CLASH algorithm," *Technical Report*, 2011.
20. Y. Censor and S. A. Zenios, *Parallel Optimization: Theory, Algorithms, and Applications*. Oxford University Press, 1997.
21. L. M. Bregman, "The relaxation method of finding the common point of convex sets and its application to the solution of problems in convex programming," *USSR Computational Mathematics and Mathematical Physics*, vol. 7, no. 3, pp. 200–217, 1967.
22. N. Cesa-Bianchi and G. Lugosi, *Prediction, Learning, and Games*. Cambridge University Press, 2006.
23. R. Berinde, A. C. Gilbert, P. Indyk, H. Karloff, and M. J. Strauss, "Combining geometry and combinatorics: A unified approach to sparse signal recovery," *preprint*, 2008.
24. S. Jafarpour, "Deterministic Compressed Sensing," Ph.D. dissertation, Princeton University, 2011.
25. N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, *Algorithmic Game Theory*. New York: Cambridge University Press, 2007.
26. J. V. Neumann, "Zur theorie der gesellschaftsspiele," *Math. Annalen*, vol. 100, pp. 295–320, 1928.
27. Y. Freund and R. E. Schapire, "Game theory, on-line prediction and boosting," in *Proceedings of the Ninth Annual Conference on Computational Learning Theory*, 1996, pp. 325–332.
28. T. H. Kjeldsen, "John von Neumann's Conception of the Minimax Theorem: A Journey through Different Mathematical Contexts," *Archive for History of Exact Sciences*, vol. 56, pp. 39–68, 2001.
29. Y. Freund and R. E. Schapire, "Adaptive game playing using multiplicative weights," *Games and Economic Behavior*, vol. 29, pp. 79–103, 1999.
30. Y. Nesterov, "Smooth minimization of non-smooth functions," *Mathematical Programming*, vol. 103, no. 1, pp. 127–152, 2005.
31. —, *Introductory lectures on convex optimization: A basic course*. Springer, 2004.

32. E. Hazan, "A survey: The convex optimization approach to regret minimization," 2011, preprint available at <http://ie.technion.ac.il/~ehazan/papers/OCO-survey.pdf>.
33. A. J. Grove, N. Littlestone, and D. Schuurmans, "General convergence results for linear discriminant updates," *Machine Learning*, vol. 43, p. 173210, 2001.
34. J. Kivinen and M. K. Warmuth, "Relative loss bounds for multi-dimensional regression problems," *Machine Learning*, vol. 45, no. 3, pp. 301–329, 2001.
35. A. Nemirovski and D. Yudin, *Problem Complexity and Method Efficiency in Optimization*. New York: Wiley, 1983.
36. A. Beck and M. Teboulle, "Mirror descent and nonlinear projected subgradient methods for convex optimization," *Operations Research Letters*, vol. 31, pp. 167–175, 2003.
37. T. Cover, "Universal portfolios," *Mathematical Finance*, vol. 1, no. 1, pp. 1–19, 1991.
38. E. Hazan, A. Agarwal, and S. Kale, "Logarithmic regret algorithms for online convex optimization," *Machine Learning*, vol. 69, no. 2-3, pp. 169–192, 2007.
39. J. Abernethy, E. Hazan, and A. Rakhlin, "Competing in the dark: An efficient algorithm for bandit linear optimization," in *the 21st Annual Conference on Learning Theory (COLT)*, 2008, pp. 263–274.
40. J. D. Lafferty, S. D. Pietra, and V. D. Pietra, "Statistical learning algorithms based on Bregman distances," in *Proceedings of the Canadian Workshop on Information Theory*, 1997.
41. M. Collins, R. E. Schapire, and Y. Singer, "Logistic regression, AdaBoost and Bregman distances," *Machine Learning*, vol. 48, no. 1/2/3, 2002.
42. S. Foucart, "Sparse recovery algorithms: sufficient conditions in terms of restricted isometry constants," in *Proceedings of the 13th International Conference on Approximation Theory*, 2010.
43. E. J. Candès and T. Tao, "Rejoinder: the Dantzig selector: statistical estimation when p is much larger than n ," *Annals of Statistics*, vol. 35, pp. 2392–2404, 2007.
44. D. L. Donoho and Y. Tsaig, "Fast solution of ℓ_1 -norm minimization problems when the solution may be sparse," *IEEE Transactions on Information Theory*, vol. 54, no. 11, pp. 4789–4812, 2008.
45. E. J. Candès and J. K. Romberg, "Quantitative robust uncertainty principles and optimally sparse decompositions," *Foundations of Computational Mathematics*, vol. 6, pp. 227–254, 2004.
46. M. S. Asif and J. K. Romberg, "Dantzig selector homotopy with dynamic measurements," in *Proceedings of Computational Imaging*, 2009.
47. E. van den Berg and M. P. Friedlander, "Probing the pareto frontier for basis pursuit solutions," *SIAM Journal on Scientific Computing*, 2008.